

Subscribe (Full Service) Register (Limited Service, Free) Login

Search:
The ACM Digital Library
The Guide

+distributed +dns

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Published before November 2001 Terms used distributed dns

Found **540** of **123,999**

Sort results

Display results

relevance expanded form

Save results to a Binder ? Search Tips Open results in a new window

Try an Advanced Search Try this search in The ACM Guide

Results 1 - 20 of 200

Result page: **1** <u>2</u> <u>3</u> <u>4</u> <u>5</u> <u>6</u> <u>7</u> <u>8</u> <u>9</u> <u>10</u>

Relevance scale 🔲 📟 📟 📟

Best 200 shown

Application level performance: DNS performance and the effectiveness of caching

Jaeyeon Jung, Emil Sit, Hari Balakrishnan, Robert Morris

November 2001 Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement

Publisher: ACM Press

Full text available: pdf(2.84 MB)

Additional Information: full citation, abstract, references, citings, index

This paper presents a detailed analysis of traces of DNS and associated TCP traffic collected on the Internet links of the MIT Laboratory for Computer Science and the Korea Advanced Institute of Science and Technology (KAIST). The first part of the analysis details how clients at these institutions interact with the wide-area DNS system, focusing on performance and prevalence of failures. The second part evaluates the effectiveness of DNS caching.In the most recent MIT trace, 23% of lookups rece ...

2. Load distribution among replicated Web servers: a QoS-based approach

Marco Conti, Enrico Gregori, Fabio Panzieri

March 2000 ACM SIGMETRICS Performance Evaluation Review, Volume 27 Issue 4

Publisher: ACM Press

Full text available: Policy pdf(695.59 KB) Additional Information: full citation, abstract, index terms

A dominant factor for the success of an Internet based Web service is the Quality of Service (QoS) perceived by its users. The principal QoS attributes these users perceive include those related to the service "responsiveness", i.e. the service availability and timeliness. In this paper, we argue that QoS can be provided by distributing the processing load among replicated Web servers, and that these servers can be geographically distributed across the Internet. In this context, we discuss strat ...

Keywords: QoS, Web server, load distribution

3 Application level performance: On the use and performance of content distribution



4/28/06

networks

ď

Balachander Krishnamurthy, Craig Wills, Yin Zhang

November 2001 Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement

Publisher: ACM Press

Full text available: pdf(2.51 MB)

Additional Information: full citation, abstract, references, citings, index

Content distribution networks (CDNs) are a mechanism to deliver content to end users on behalf of origin Web sites. Content distribution offloads work from origin servers by

serving some or all of the contents of Web pages. We found an order of magnitude increase in the number and percentage of popular origin sites using CDNs between November 1999 and December 2000.In this paper we discuss how CDNs are commonly used on the Web and define a methodology to study how well they perform. A performanc ...

4 <u>Dynamically distributed query evaluation</u>

Trevor Jim, Dan Suciu

May 2001 Proceedings of the twentieth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems

Publisher: ACM Press

Full text available: pdf(276.64 KB) Additional Information: full citation, references, citings, index terms

⁵ Passive measurements: Characterizing large DNS traces using graphs

Charles D. Cranor, Emden Gansner, Balachander Krishnamurthy, Oliver Spatscheck November 2001 Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement

Publisher: ACM Press

۵

Full text available: pdf(2.62 MB)

Additional Information: full citation, abstract, references, citings, index terms

The increasing deployment of overlay networks that rely on DNS tricks has led to added interest in examining DNS traffic. In this paper we report on a characterization of DNS traffic gathered over a period of several weeks at Internet Gateway Routers (IGRs) in the AT&T Common Backbone. The characterization is carried out using several novel techniques to identify clients, local DNS servers, and authoritative DNS servers. Our techniques include passive and active measurements, graph-based analysi ...

⁶ A distributed hypercube file system

, R. J. Flynn, H. Hadimioglu

January 1989 Proceedings of the third conference on Hypercube concurrent computers and applications - Volume 2

Publisher: ACM Press

Full text available: pdf(651.48 KB)

Additional Information: full citation, abstract, references, citings, index terms

For the hypercube, an autonomous physically interconnected file system is proposed. The resulting distributed file system consists of an I/O organization and a software interface. The system is loosely-coupled architecturally but from operating systems point of view a tightly-coupled system is formed in which interprocessor messages are handled differently from file accesses. A matrix multiplication algorithm is given to show how the distributed file system is utilized.

7. Adaptive TTL schemes for load balancing of distributed Web servers

Michele Colajanni, Philip S. Yu

September 1997 ACM SIGMETRICS Performance Evaluation Review, Volume 25 Issue 2

Publisher: ACM Press

Full text available: pdf(488.09 KB) Additional Information: full citation, abstract, citings, index terms

With ever increasing web traffic, a distributed Web system can provide scalability and flexibility to cope with growing client demands. Load balancing algorithms to spread the load across multiple Web servers are crucial to achieve the scalability. Various *domain name server* (DNS) based schedulers have been proposed in the literature, mainly for multiple homogeneous servers. DNS provides (logical) host name to IP-address mapping (i.e., the server assignment), but the mapping is not done f ...

8 <u>Distributed cooperative Apache web server</u>

Quanzhong Li, Bongki Moon

April 2001 Proceedings of the 10th international conference on World Wide Web

Publisher: ACM Press

Full text available: pdf(306.36 KB) Additional Information: full citation, references, citings, index terms

Keywords: Apache, DC-Apache, WWW, distributed Web server, load balancing, replication, scalable Web server

9 Development of the Domain Name System

Paul V. Mockapetris, Kevin J. Dunlap

January 1995 ACM SIGCOMM Computer Communication Review, Volume 25 Issue 1

Publisher: ACM Press

į,

Full text available: pdf(983.50 KB) Additional Information: full citation, abstract, citings, index terms

The Domain Name System (DNS) provides name service for the DARPA Internet. It is one of the largest name services in operation today, serves a highly diverse community of hosts, users, and networks, and uses a unique combination of hierarchies, caching, and datagram access. This paper examines the ideas behind the initial design of the DNS in 1983, discusses the evolution of these ideas into the current implementations and usages, notes conspicuous surprises, successes and shortcomings, and attem ...

10 Design and evaluation of a replicated database for mobile systems

S. Palazzo, A. Puliafito, M. Scarpa

March 2000 Wireless Networks, Volume 6 Issue 2

Publisher: Kluwer Academic Publishers

Full text available: pdf(251.84 KB) Additional Information: full citation, abstract, references, index terms

The new generation mobile systems are anticipated to provide mobile users with new broadband services such as wireless multimedia, including real‐ time video and high‐ speed data. In these systems, the requirement on service transparency placed by the handling of mobility, both personal and terminal, implies a remarkable increase in the complexity of data management. Therefore, appropriate distributed databases (DDB) must be designed to guarantee speed in the processing an ...

11 Development of the domain name system

P. Mockapetris, K. J. Dunlap

August 1988 ACM SIGCOMM Computer Communication Review , Symposium proceedings on Communications architectures and protocols SIGCOMM

'88, Volume 18 Issue 4

Publisher: ACM Press

Full text available: pdf(1.24 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u>

The Domain Name System (DNS) provides name service for the DARPA Internet. It is one of the largest name services in operation today, serves a highly diverse community of hosts, users, and networks, and uses a unique combination of hierarchies, caching, and datagram access. This paper examines the ideas behind the initial design of the DNS in 1983, discusses the evolution of these ideas into the current implementations and usages, notes conspicuous surprises, successes and shortc ...

12 Timing attacks on Web privacy

Edward W. Felten, Michael A. Schneider

November 2000 Proceedings of the 7th ACM conference on Computer and communications security

Publisher: ACM Press

Full text available: pdf(184.79 KB) Additional Information: full citation, references, citings, index terms

13 Scheduling optimization for resource-intensive Web requests on server clusters

Huican Zhu, Ben Smith, Tao Yang

Unne 1999 Proceedings of the eleventh annual ACM symposium on Parallel algorithms and architectures



Publisher: ACM Press

Full text available: pdf(1.19 MB) Additional Information: full citation, references, citings, index terms

14 An analysis of wide-area name server traffic: a study of the Internet Domain Name



<u>System</u>

Peter B. Danzig, Katia Obraczka, Anant Kumar

October 1992 ACM SIGCOMM Computer Communication Review, Conference proceedings on Communications architectures & protocols SIGCOMM

'92, Volume 22 Issue 4

Publisher: ACM Press

Full text available: pdf(1.19 MB)

Additional Information: full citation, abstract, references, citings, index terms

Over a million computers implement the Internet's Domain Name System of DNS, making it the world's most distributed database and the Internet's most significant source of wide-area RPC-like traffic. Last year, over eight percent of the packets and four percent of the bytes that traversed the NSFnet were due to DNS. We estimate that a third of this wide-area DNS traffic was destined to seven root name servers. This paper explores the performance of DNS based on two 24-hour t ...

15 Dynamic Load Balancing DNS

V. C. Harish, Brad Owens

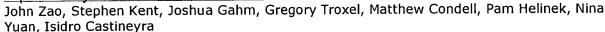
August 1999 Linux Journal

Publisher: Specialized Systems Consultants, Inc.

Full text available: html(12.52 KB) Additional Information: full citation, abstract, references, index terms

dlbDNS: This article discusses an attempt to solve the problem of network traffic congestionadding a dynamic load balancing feature to the existing DNS

16 A public-key based secure mobile IP



September 1997 Proceedings of the 3rd annual ACM/IEEE international conference on Mobile computing and networking

Publisher: ACM Press

Additional Information: full citation, references, citings Full text available: pdf(1.95 MB)

17 Distributed indexing: a scalable mechanism for distributed information retrieval



Peter B. Danzig, Jongsuk Ahn, John Noll, Katia Obraczka

September 1991 Proceedings of the 14th annual international ACM SIGIR conference on Research and development in information retrieval

Publisher: ACM Press

Full text available: 📆 pdf(929.40 KB) Additional Information: full citation, references, citings, index terms

18 Take Command: An Introduction to DNS and DNS Tools

Neil Anuskiewicz

May 2001 Linux Journal

Publisher: Specialized Systems Consultants, Inc.

Full text available: html(11.02 KB) Additional Information: full citation, index terms

19 A scalable content-addressable network



Sylvia Ratnasamy, Paul Francis, Mark Handley, Richard Karp, Scott Schenker

August 2001 ACM SIGCOMM Computer Communication Review, Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '01, Volume 31 Issue 4

Publisher: ACM Press

Full text available: pdf(155.64 KB)

Additional Information: full citation, abstract, references, citings, index terms

Hash tables - which map "keys" onto "values" - are an essential building block in modern software systems. We believe a similar functionality would be equally valuable to large distributed systems. In this paper, we introduce the concept of a Content-Addressable Network (CAN) as a distributed infrastructure that provides hash table-like functionality on Internet-like scales. The CAN is scalable, fault-tolerant and completely self-organizing, and we demonstrate its scalability, robustness and low ...

20 Applying an information gathering architecture to Netfind: a white pages tool for a



<u>changing and growing Internet</u> Michael F. Schwartz, Calton Pu

October 1994 IEEE/ACM Transactions on Networking (TON), Volume 2 Issue 5

Publisher: IEEE Press

Full text available: pdf(1.71 MB)

Additional Information: <u>full citation</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>, <u>review</u>

Results 1 - 20 of 200

ذ

Result page: 1 2 3 4 5 6 7 8 9 10 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Mindows Media Player Real Player



Subscribe (Full Service) Register (Limited Service, Free) Login

Search:

The ACM Digital Library
The Guide

peer-to-peer dns

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used peer to peer dns

Found 2,068 of 175,083

Sort results relevance by Display expanded form

Save results to a Binder Search Tips Open results in a new

Try an Advanced Search Try this search in The ACM Guide

Results 1 - 20 of 200

window

Result page: 1 2 3 4 5 6 7 8 9 10

Relevance scale 🔲 📟 📟 👹

Best 200 shown

reșults

À

٤

DNS: Availability, usage, and deployment characteristics of the domain name system

Jeffrey Pang, James Hendricks, Aditya Akella, Roberto De Prisco, Bruce Maggs, Srinivasan

October 2004 Proceedings of the 4th ACM SIGCOMM conference on Internet measurement

Publisher: ACM Press

Full text available: pdf(856.34 KB) Additional Information: full citation, abstract, references, index terms

The Domain Name System (DNS) is a critical part of the Internet's infrastructure, and is one of the few examples of a robust, highly-scalable, and operational distributed system. Although a few studies have been devoted to characterizing its properties, such as its workload and the stability of the top-level servers, many key components of DNS have not yet been examined. Based on large-scale measurements taken fromservers in a large content distribution network, we present a detailed study of ...

Keywords: DNS, availability, federated

Chord: a scalable peer-to-peer lookup protocol for internet applications Ion Stoica, Robert Morris, David Liben-Nowell, David R. Karger, M. Frans Kaashoek, Frank Dabek, Hari Balakrishnan

February 2003 IEEE/ACM Transactions on Networking (TON), Volume 11 Issue 1

Publisher: IEEE Press

Full text available: pdf(690.54 KB)

Additional Information: full citation, abstract, references, citings, index <u>terms</u>

A fundamental problem that confronts peer-to-peer applications is the efficient location of the node that stores a desired data item. This paper presents Chord, a distributed lookup protocol that addresses this problem. Chord provides support for just one operation: given a key, it maps the key onto a node. Data location can be easily implemented on top of Chord by associating a key with each data item, and storing the key/data pair at the node to which the key maps. Chord adapts efficien ...

Keywords: distributed scalable algorithms, lookup protocols, peer-to-peer networks

The design and implementation of a next generation name service for the internet Venugopalan Ramasubramanian, Emin Gün Sirer

August 2004 ACM SIGCOMM Computer Communication Review, Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '04, Volume 34 Issue 4

Publisher: ACM Press



Full text available: pdf(472.93 KB) Additional Information: full citation, abstract, references, citings, index

Name services are critical for mapping logical resource names to physical resources in large-scale distributed systems. The Domain Name System (DNS) used on the Internet, however, is slow, vulnerable to denial of service attacks, and does not support fast updates. These problems stem fundamentally from the structure of the legacy DNS. This paper describes the design and implementation of the Cooperative Domain Name System (CoDoNS), a novel name service, which provides high lookup performance thro ...

Keywords: DNS, peer to peer, proactive caching

4. Article abstracts with full text online: The liquid architecture: a non-linear peer-to-peer distributed architecture with polymorphic message passing

Coskun Bayrak, Chad Davis

May 2003 ACM SIGSOFT Software Engineering Notes, Volume 28 Issue 3

Publisher: ACM Press

Full text available: pdf(296.94 KB) Additional Information: full citation, abstract, references

In terms of benefiting from the potential to be gained from full distribution, today's most common implementations of distributed systems follow only limited linear versions of distribution such as client-server or n-tier models. Even many "peer to peer" systems still rely on centralized servers to provide the message passing connectivity between the peers. While these systems do provide increased robustness and computational speedup, they fail to realize the full measure of what fully distribut ...

Keywords: P2P, architecture, distributed systems, networking, peer-to-peer, real time systems, software engineering, virtual collaboration

5; A scalable distributed information management system



Praveen Yalagandula, Mike Dahlin

August 2004 ACM SIGCOMM Computer Communication Review, Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '04, Volume 34 Issue 4

Publisher: ACM Press

Full text available: pdf(364.00 KB) Additional Information: full citation, abstract, references, index terms

We present a Scalable Distributed Information Management System (SDIMS) that aggregates information about large-scale networked systems and that can serve as a basic building block for a broad range of large-scale distributed applications by providing detailed views of nearby information and summary views of global information. To serve as a basic building block, a SDIMS should have four properties: scalability to many nodes and attributes, flexibility to accommodate a broad range of appl ...

Keywords: distributed hash tables, information management system, networked system monitoring

IPNL: A NAT-extended internet architecture



Paul Francis Ramakrishna

August 2001 ACM SIGCOMM Computer Communication Review, Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '01, Volume 31 Issue 4

Publisher: ACM Press

Full text available: pdf(241.65 KB)

Additional Information: full citation, abstract, references, citings, index

This paper presents and analyzes IPNL (for IP Next Layer), a NAT-extended Internet protocol architecture designed to scalably solve the address depletion problem of IPv4. A NAT-extended architecture is one where only hosts and NAT boxes are modified. IPv4 routers and support protocols remain untouched. IPNL attempts to maintain all of the original characteristics of IPv4, most notably address prefix location independence. IPNL provides true site isolation (no renumbering), and allows sites to be ...

Measuring the internet's vital statistics: Observations of IPv6 traffic on a 6to4 relay

Pekka Savola

January 2005 ACM SIGCOMM Computer Communication Review, Volume 35 Issue 1



Publisher: ACM Press

Full text available: pdf(186.86 KB) Additional Information: full citation, abstract, references, index terms

FUNET has been operating a public, globally-used 6to4 (RFC 3056) relay router since November 2001. The traffic has been logged and is now analyzed to gather information of 6to4 and IPv6 deployment. Among other figures, we note that the number of 6to4 capable nodes has increased by an order of magnitude in half a year: in April 2004, there are records of about 2 million different 6to4 nodes using this particular relay. Vast majority of this is just testing the availability of the relay, done by th ...

Keywords: 6to4, IPv6, IPv6 transition

& Chord: A scalable peer-to-peer lookup service for internet applications

Ion Stoica, Robert Morris, David Karger, M. Frans Kaashoek, Hari Balakrishnan
August 2001 ACM SIGCOMM Computer Communication Review, Proceedings of the
2001 conference on Applications, technologies, architectures, and
protocols for computer communications SIGCOMM '01, Volume 31 Issue 4

Publisher: ACM Press

Full text available: pdf(205.73 KB)

Additional Information: full citation, abstract, references, citings, index terms

A fundamental problem that confronts peer-to-peer applications is to efficiently locate the node that stores a particular data item. This paper presents *Chord*, a distributed lookup protocol that addresses this problem. Chord provides support for just one operation: given a key, it maps the key onto a node. Data location can be easily implemented on top of Chord by associating a key with each data item, and storing the key/data item pair at the node to which the key maps. Chord adapts effi ...

9 Incrementally improving lookup latency in distributed hash table systems

Hui Zhang, Ashish Goel, Ramesh Govindan

June 2003 ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2003 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '03, Volume 31 Issue 1

Publisher: ACM Press

Full text available: pdf(401.28 KB)

Additional Information: full citation, abstract, references, citings, index terms

Distributed hash table (DHT) systems are an important class of peer-to-peer routing infrastructures. They enable scalable wide-area storage and retrieval of information, and will support the rapid development of a wide variety of Internet-scale applications ranging from naming systems and file systems to application-layer multicast. DHT systems essentially build an overlay network, but a path on the overlay between any two nodes can be significantly different from the unicast path between those ...

Keywords: DHT, latency stretch, peer-to-peer, random sampling

PocketLens: Toward a personal recommender system
Bradley N. Miller, Joseph A. Konstan, John Riedl

July 2004 ACM Transactions on Information Systems (TOIS), Volume 22 Issue 3

Publisher: ACM Press



٤

Full text available: pdf(1.10 MB)

Additional Information: full citation, abstract, references, index terms

Recommender systems using collaborative filtering are a popular technique for reducing information overload and finding products to purchase. One limitation of current recommenders is that they are not portable. They can only run on large computers connected to the Internet. A second limitation is that they require the user to trust the owner of the recommender with personal preference data. Personal recommenders hold the promise of delivering high quality recommendations on palmtop computers, e ...

Keywords: Collaborative Filtering, Peer-to-Peer Networking, Privacy, Recommender **Systems**

Panel Session: Peer to peer: still useless?

October 2005 Proceedings of the twentieth ACM symposium on Operating systems principles SOSP '05

Publisher: ACM Press

Full text available: pdf(83.50 KB) Additional Information: full citation

13 A churn-resistant peer-to-peer web caching system

Prakash Linga, Indranil Gupta, Ken Birman

October 2003 Proceedings of the 2003 ACM workshop on Survivable and selfregenerative systems: in association with 10th ACM Conference on **Computer and Communications Security**

Publisher: ACM Press

Full text available: pdf(1.07 MB)

Additional Information: full citation, abstract, references

Denial of service attacks on peer-to-peer (p2p) systems can arise from sources otherwise considered non-malicious. We focus on one such commonly prevalent source, called "churn". Churn arises from continued and rapid arrival and failure (or departure) of a large number of participants in the system, and traces from deployments have shown that it can lead to extremely stressful networking conditions. It has the potential to increase host loads and block a large fraction of normal insert and lo ...

13 SelectCast: a scalable and self-repairing multicast overlay routing facility

Adrian Bozdog, Robbert van Renesse, Dan Dumitriu

October 2003 Proceedings of the 2003 ACM workshop on Survivable and selfregenerative systems: in association with 10th ACM Conference on **Computer and Communications Security**

Publisher: ACM Press

Full text available: pdf(1.01 MB)

Additional Information: full citation, abstract, references

In this paper we describe SelectCast, a self-repairing multicast overlay routing facility for supporting publish/subscribe applications. Select Cast is a peer-to-peer protocol, and lever-ages Astrolabe, a secure distributed information management system. SelectCast uses replication to recover quickly from transient failures, as well as Astrolabe's aggregation facilities to recover from long-term failures or adapt to changes in load or QoS requirements. In order to evaluate the scalability and ...

14 Improving lookup latency in distributed hash table systems using random sampling Hui Zhang, Ashish Goel, Ramesh Govindan

October 2005 IEEE/ACM Transactions on Networking (TON), Volume 13 Issue 5

Publisher: IEEE Press

٤

Full text available: pdf(511.97 KB) Additional Information: full citation, abstract, references, index terms

Distributed hash table (DHT) systems are an important class of peer-to-peer routing infrastructures. They enable scalable wide-area storage and retrieval of information, and will support the rapid development of a wide variety of Internet-scale applications ranging from naming systems and file systems to application-layer multicast. DHT systems

à

essentially build an overlay network, but a path on the overlay between any two nodes can be significantly different from the unicast path between those ...

Keywords: distributed hash table (DHT), internet topology, latency expansion, latency stretch, peer-to-peer, random sampling, randomized algorithm

15 Network behavior: An analysis of Internet content delivery systems



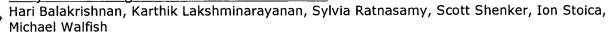
Stefan Saroiu, Krishna P. Gummadi, Richard J. Dunn, Steven D. Gribble, Henry M. Levy December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI

· Publisher: ACM Press

Full text available: pdf(2.07 MB) Additional Information: full citation, abstract, references

In the span of only a few years, the Internet has experienced an astronomical increase in the use of specialized content delivery systems, such as content delivery networks and peer-to-peer file sharing systems. Therefore, an understanding of content delivery on the Internet now requires a detailed understanding of how these systems are used in practice. This paper examines content delivery from the point of view of four content delivery systems: HTTP web traffic, the Akamai content delivery netw ...

16 A layered naming architecture for the internet



August 2004 ACM SIGCOMM Computer Communication Review , Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '04, Volume 34 Issue 4

Publisher: ACM Press

Full text available: pdf(110.95 KB)

Additional Information: full citation, abstract, references, citings, index terms

Currently the Internet has only one level of name resolution, DNS, which converts user-level domain names into IP addresses. In this paper we borrow liberally from the literature to argue that there should be three levels of name resolution: from user-level descriptors to service identifiers; from service identifiers to endpoint identifiers; and from endpoint identifiers to IP addresses. These additional levels of naming and resolution (1) allow services and data to be first class Internet objec ...

Keywords: distributed hash tables, global identifiers, internet architecture, middleboxes, name resolution, naming

17 Distributed hash tables, Part I

Brandon Wiley

October 2003 Linux Journal, Volume 2003 Issue 114

Publisher: Specialized Systems Consultants, Inc.

Full text available: html(19.41 KB) Additional Information: full citation, abstract

Learn the fundamental technique behind the next generation of privacy-conscious peer-to-peer systems.

18 Systems and applications: DAIMON: data integration for a mobile network



June 2005 Proceedings of the 4th ACM international workshop on Data engineering for wireless and mobile access

Publisher: ACM Press

Full text available: pdf(365.53 KB) Additional Information: full citation, abstract, references, index terms

We describe the DAIMON system for data integration of nomadic data. DAIMON is based on an extension of the mediator-wrapper architecture, where each node is a mediator of its own information and information from its neighbouring nodes. We describe the overall architecture, including our lightweight peer location and neighbourhood identification

scheme, cache management strategy and query processing and execution architecture, and show how it addresses the challenges of the wireless environment. ...

Keywords: connectivity profile, mediator, mobile data management

19 Features: Designing Portable Collaborative Networks

May 2003 Queue, Volume 1 Issue 3

Publisher: ACM Press

Full text available: pdf(547.14 KB) Additional Information: full citation, index terms html(41.58 KB)

Session 6: Viceroy: a scalable and dynamic emulation of the butterfly



Dahlia Malkhi, Moni Naor, David Ratajczak

July 2002 Proceedings of the twenty-first annual symposium on Principles of distributed computing

Publisher: ACM Press

Full text available: Pdf(1.09 MB) Additional Information: full citation, abstract, references, citings

We propose a family of constant-degree routing networks of logarithmic diameter, with the additional property that the addition or removal of a node to the network requires no global coordination, only a constant number of linkage changes in expectation, and a logarithmic number with high probability. Our randomized construction improves upon existing solutions, such as balanced search trees, by ensuring that the congestion of the network is always within a logarithmic factor of the optimum with ...

Results 1 - 20 of 200 Result page: 1 2 3 4 5 6 7 8 9 10

The ACM Portal is published by the Association for Computing Machinery. Copyright @ 2006 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Mindows Media Player

EAST Search History

Ref	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1 ^t	12026287	@ad<"20011113"	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2006/04/28 16:24
L2	11855	common with zon\$5	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2006/04/28 16:24
L3.	36935	lookup with tabl\$5	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2006/04/28 16:25
L4	32523	dns or (domain adj name adj servic\$5)	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2006/04/28 16:26
L5	387	name with resolution with protocol\$5	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2006/04/28 16:26
L6	1	2 and 3 and 4 and 5	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2006/04/28 16:26
L7	0	1 and 6	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2006/04/28 16:26
L8	58	2 and 4 and 1	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2006/04/28 16:27
L9 ⁻	9895	address with resolution	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2006/04/28 16:27
L10	1129	peer and 9	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2006/04/28 16:27
L11	526	1 and 10	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2006/04/28 16:27

4/28/06 4:28:40 PM Page 1

EAST Search History

L12	11	2 and 11	US-PGPUB;	OR	ON	2006/04/28 16:27
			USPAT;			
160			EPO; JPO;			
	'		IBM_TDB			

4/28/06 4:28:40 PM Page 2